CISTO news

Fall/Winter 2005

Code 606 • http://cisto.gsfc.nasa.gov
Computational and Information Sciences and Technology Office
Sciences and Exploration Directorate, Goddard Space Flight Center

In This Issue

CISTO Update

NASA Makes Strong Impression at Record-Setting SC05 Conference

A record 9,777 people converged on Seattle, WA, November 12–18 for SC05, the premiere international conference for high-performance computing, networking, storage, and analysis. This participation represented a 22% increase over 2004. Attendees from school-age children to seasoned scientists visited NASA's engaging research exhibit, which had a high-profile location adjacent to Intel Corp.

The NASA exhibit showcased a variety of applications from all four Mission Directorates: Aeronautics, Exploration Systems, Science, and Space Operations. Demonstrations were on workstation pedestals along the exhibit perimeter and in a central

NASA's SC05 research exhibit featured applications from all four Mission Directorates. Goddard Space Flight Center contributed six demonstrations on Earth and space science topics (Photo credit: David Robertson, NASA Ames Research Center).

presentation area outfitted with a large plasma screen and a nine-screen HyperWall. Researchers connected to Goddard Space Flight Center (GSFC) presented six demonstrations on Earth and space science topics.

Gaining deeper insights from Earth observations was the theme of "Interactive Image Segmentation with RHSEG and HSEGViewer." Jim Tilton of CISTO showed these software tools for producing and then visualizing and manipulating image segmentation hierarchies, which contain several segmentations, from coarser to higher levels of detail, of the same image. Among the capabilities are classification and interactive labeling of observation images. Tilton's software is also being used for medical and other applications (see "CISTO Engineer Receives Patent and GSFC IS&T Award," *CISTO News*, Summer 2005).

Two demonstrations detailed recent advances made using Computational Tech-

nologies Project-funded software frameworks.

In "Cross-Organization Coupling of Climate Models through ESMF," Shujia Zhou of the Software Integration and Visualization Office (SIVO)/Northrop Grumman IT described how the Earth System Modeling Framework (ESMF) has enabled coupling of models from seven organizations within a standardized software environment. Zhou also explained an updated prototype based on ESMF and Common Component Architecture software. In this prototype, the representative atmosphere and ocean models ran on two separate Thunderhead cluster partitions connected to the CISTO High End Computer Network (HECN) team's 10-gigabit-per-second (Gbps) Lambda Network, with model data being exchanged over the regional 10-Gbps DRAGON network.

In "Columbia Enables New Space Weather Modeling," Darren De Zeeuw of the University of Michigan showed Space Weather Modeling Framework (SWMF) simulations of the Halloween 2003 space storms run on the Columbia supercomputer at NASA Ames Research Center. The simulations depict complex magnetic field configurations in the sun's corona as well as the response of Earth's magnetosphere to some of the most powerful solar eruptions to ever shake the solar system. Complementing the demonstration, SWMF wall posters were available for attendees to take home (see "New Space Weather Poster Available," *CISTO News*, Summer 2005).

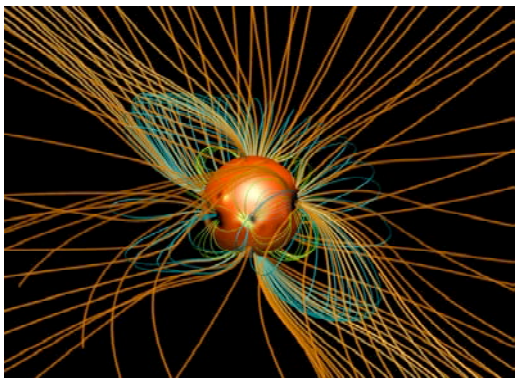

A Space Weather Modeling Framework simulation of the Halloween 2003 space storms depicts the structure of the large-scale solar magnetic fields on October 27, 2003. (Image credit: Ilia Roussev, University of Michigan).

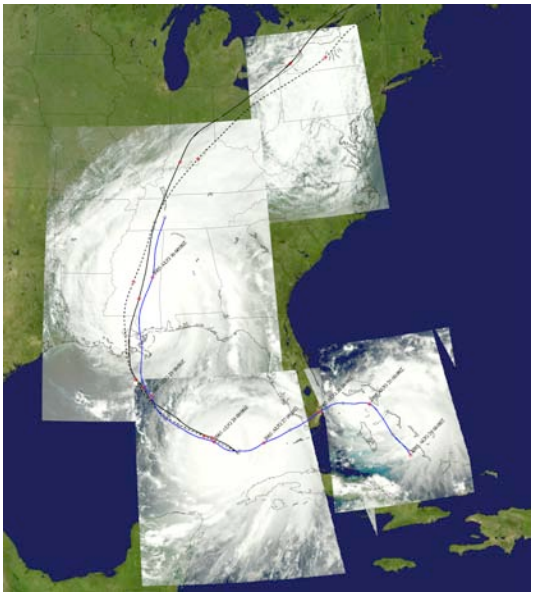

Hurricane Katrina's observed track is overlaid in blue along with the predicted tracks from the GEOS4 (solid black) and GEOS5 (dashed black) models initialized at 12:00Z on August 27, 2005. The cloud imagery comes from the MODIS sensor on NASA's Aqua satellite (Image credit: Bill Putman, GSFC).

The unprecedented Fall 2005 Atlantic hurricane season was the subject of two demonstrations by GSFC scientists. In "MAP '05—Project Hurricane," Bill Putman of SIVO described how GSFC's GEOS4 and GEOS5 Atmospheric General Circulation Models were run on Columbia at .25-degree resolution four times per day throughout the hurricane season. Through inclusion of storm tracks and other output in Florida State University's "Superensemble," these global models provided real-time guidance to National Hurricane Center forecasters.

Dan Kokron of the Global Modeling and Assimilation Office/SAIC presented "GEOS5—Columbia and Hurricanes." GEOS5 is the first major operational system built from the ground up using ESMF and its

object-oriented concepts. Notably, GEOS5 was one of the earliest models to predict a New Orleans landfall for Hurricane Katrina and was consistent in its landfall prediction from two days beforehand.

Hurricane modeling presentations on the HyperWall included live feeds of data sets exchanged between Columbia and CISTO's NASA Center for Computational Sciences over the recently upgraded 1-Gbps NASA Research and Engineering Network path across the National LambdaRail.

General relativity came to the fore in "Modeling Gravitational Wave Sources for LISA," presented by the Gravitational Astrophysics Laboratory's Dae-Il Choi (USRA), Michael Koppitz (NRC), and Jim Van Meter (NRC). Using Columbia, the group has been modeling mergers of two comparable-mass black holes and calculating the resulting gravitational wave signatures. The calculated waveforms will be applied to analyzing and interpreting observations from the Laser Interferometer Space Antenna (LISA), a joint NASA-European Space Agency mission.

As GSFC coordinator, Jarrett Cohen of CISTO/GST, Inc. participated in NASA exhibit planning and preparation and produced publicity materials for the demonstrations.

Across the show floor in the Internet2 exhibit, the HECN team supported a demonstration of the electronic-Very Long Baseline Interferometry (e-VLBI) project. Employing the same set of technologies used for iGRID 2005 (see "CISTO Supports iGrid 2005 Demonstrations" in this issue), the demonstration correlated radio telescope data from e-VLBI sites in the United States, Sweden, and Japan in real time.

http://sc05.supercomputing.org/

Networks & IT Security

CISTO Supports iGrid 2005 Demonstrations

Two live demonstrations went off without a hitch at the iGrid 2005 Workshop, with support from CISTO's High End Computer Network (HECN) team and their 10 gigabit-per-second fiber optic network. Using the DRAGON regional optical network and the National LambdaRail (NLR), data and video flowed 3,000 miles between Goddard Space Flight Center (GSFC) and the workshop site at the University of California, San Diego.

An iGrid 2005 demonstration of the electronic-Very Long Baseline Interferometry (e-VLBI) project correlated data from telescopes worldwide (Image credit: MIT Haystack Observatory).

The California Institute for Telecommunications and Information Technology (CalIT2) hosted iGrid 2005 to encourage the development and use of high-speed optical networks. Researchers from 20 countries staged 49 exhibits during the workshop, which drew nearly 450 attendees September 26–29.

One exhibit featured the electronic-Very Long Baseline Interferometry (e-VLBI) project, a multi-national effort involving the GSFC Geophysical and Astronomical Observatory (GGAO) and the Massachusetts Institute of Technology (MIT). The demonstration used a technology called dynamic provisioning to automatically establish network connections between telescopes in Japan, the Netherlands, Sweden, and the United Kingdom. Live data from these telescopes were delivered at over 512 megabits per second to MIT, where they were correlated with data from GGAO in real time.

A second exhibit demonstrated a novel high-definition 3D video display system. Standing in front of a 4-foot-tall display podium, attendees could view holographic 3D video at HDTV resolution—without goggles or special headgear. The video originated from the HECN network lab at GSFC, where specially aligned HDTV cameras imaged a scene and encoded it for network transfer. The resulting data stream traveled through the DRAGON and NLR networks to the 3D display on the iGrid 2005 exhibit floor.

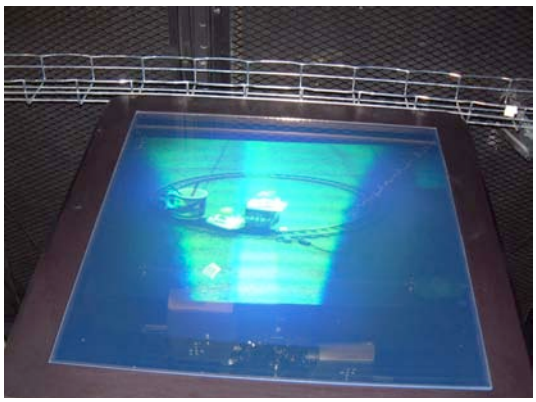

A True-3D display showed real-time, stereo-HDTV video transmitted from GSFC to the University of California, San Diego. (Image credit: Pat Gary, GSFC.)

Echoing the theme of iGrid 2005, the HECN team emphasizes that high-speed optical networking services are available today within several GSFC facilities. Projects interested in using 10-gigabit networks to advance their research are invited to contact Pat Gary, Network Projects Leader in GSFC Code 606.1, at Pat.Gary@nasa.gov.

http://www.igrid2005.org
http://cisto.gsfc.nasa.gov/IRAD_Lambda.html

High-Performance Computing

NCCS Joins DICE partnership

The NASA Center for Computational Sciences (NCCS), jointly with the U.S. Departments of Defense (DoD) and Energy (DoE), is establishing the Data Intensive Computing Environment (DICE). DICE, administered by AVETEC, will provide the necessary information technology proving ground to test solutions for data management problems associated with high-performance computing (HPC).

Scientists, engineers, and researchers use HPC to solve complex and interdisciplinary problems. One of the major performance bottlenecks is the difficulty in effectively managing the massive amount of data generated by the applications running on the HPC systems. The time to solution for many researchers is often hindered by the large number and size of files as well as the time and difficulty associated in transferring large files between systems. These issues directly affect HPC users every day and delay the critical analysis and processing of data.

As applications grow more complex and problem sizes increase, issues associated with managing large data sets are becoming even more difficult to solve. While great ad-

vances in network technologies have occurred, typical end users do not have the ability to acquire high-bandwidth network connections to facilitate data management.

The initial task of DICE is to establish a test environment at three sites – the NCCS, DoD's Aeronautical Systems Center (ASC) Major Shared Resource Center (MSRC), and DoE's Ohio Supercomputer Center. The test bed will support the identification, investigation, and development of cutting edge hardware and software solutions specifically designed to overcome problems associated with the creation and use of large data sets. Typically, these data sets will be distributed not only across geographically diverse locations but also across agencies over real-world wide area networks.

The DICE program is currently negotiating with various vendors to finalize the initial technical implementation. The three sites will have a combination of state-of-the-art Linux clusters, distributed file systems, storage devices, and a complete security infrastructure. Equipment is scheduled to start installation at the NCCS in early 2006.

Through the investments made by the three participating agencies and the vendor community, the ultimate goal of DICE is to create a self-sustaining test bed environment for the HPC community. This test bed environment will provide the necessary proving ground for viable HPC solutions with appropriate security measures to facilitate a shared and distributed data environment.

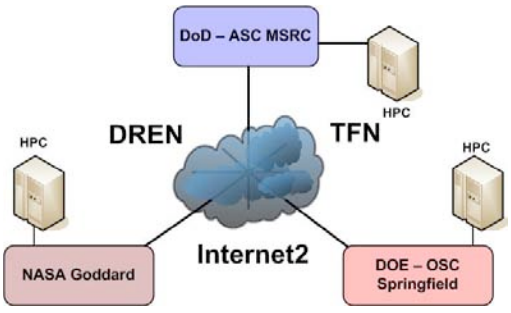

The initial DICE environment will create a collaborative testbed linking DoD, DoE, and NASA over existing networks. (Image credit: Tracey Wilson, CSC.)

http://www.avetec.org/dice
http://nccs.nasa.gov

User Allocation Process Changes

The Science Mission Directorate (SMD) has changed the face of high-performance computing (HPC) at Goddard Space Flight Center by instituting a new system of resource allocations at the NASA Center for Computational Sciences (NCCS). With this action, the SMD expects to increase overall system utilization and quantify the computing requirements of its science community.

The NCCS has teamed up with REI systems to use their e-Books software, which provides Principal Investigators (PIs) with a single repository for all documents required throughout the allocation process. The process begins with a PI formally submitting a request for resources during the given year. Because in the past it was difficult for the NCCS to predict future resource requirements for the users (e.g., procurement of additional hardware, software, or storage), a PI must answer a list of questions detailing all manner of support a computational project would require of the NCCS. Examples of information requested are computing and storage requirements, needs for third-party software, network connectivity requirements, and requests for applications support.

Once SMD determines that a PI has valid NASA funding, it begins the process of granting or denying individual allocation requests. For the initial allocation period beginning November 2005, over 100 proposals were submitted across the full range of Earth and space science disciplines. This represented a potential over-subscription of the NCCS systems available for SMD allocation. Therefore, to accommodate as many proposals as possible, many of the allocation levels were set at levels less than those requested. Because SMD intends to review

utilization throughout the life of each computational project, there is the potential for allocation levels to be adjusted quarterly.

By implementing this new process, the SMD is hoping to accomplish several goals. The first is to provide a uniform approach to allocating the shared resources at the NCCS. This will provide computing opportunity for projects that have never before run on the NCCS systems. The second goal is to increase the overall utilization of the NCCS systems. Previously, the systems were architected to accommodate a specific user community, which meant processors were often idle. By increasing the overall systems population, many users are able to run when other groups are inactive. Finally, by requiring reports from the PIs, the SMD will be better able to provide NASA Headquarters a summary of how the NCCS is supporting the individual science efforts. This in turn helps provide a justification to NASA Headquarters for additional resources as well as offering the PIs additional exposure to NASA management.

The NCCS is currently accepting proposals for its systems on a rolling basis. New proposals will be evaluated quarterly and potentially added to the overall workload on the systems. Interested parties should visit the NCCS web site.

http://www.nccs.nasa.gov/account_info.html.

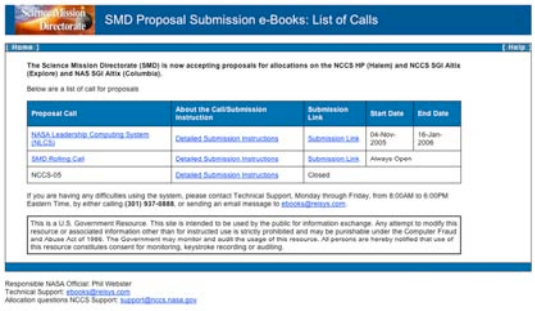

Screen shot of the Science Mission Directorate's Proposal Submission e-Books: List of Calls web page. (Image credit: Sadie Duffy, CSC.)

Information Science & Technology Research

NASA Teams with USGS to Fight Invasive Species

Each year, the United States incurs direct loses of approximately $120 billion due to non-indigenous invasive species, an annual cost greater then most natural disasters combined. Increasingly, invasive species are coming into the U.S. and spreading with globalization. An invasive species is defined as a non-native plant, animal, or microbe whose introduction causes, or is likely to cause, harm to the economy, the environment, or human health.

Previous attempts to monitor and track invasive species through field monitoring have proven too slow, costly and inefficient due, in part, to these species' rates of spread and cost of the required manpower. The U.S. Geological Survey (USGS) has partnered with NASA to use NASA Earth observations and systems engineering to enhance USGS capabilities for generating predictive models and maps of invasive species habitat and distribution. This enhanced capability within USGS, called the Invasive Species Forecasting System (ISFS), is slated to be operational in 2008. The ISFS is a computational framework that incorporates satellite data, geo-statistical models, and ecological field data to produce predictive maps that assist resource managers with invasive species management and habitat conservation.

Scientists at Goddard Space Flight Center and the USGS Fort Collins Science Center, through the Center's National Institute of Invasive Species Science in Fort Collins, CO, have been working on the ISFS project for three years. With funding from the Earth-Sun System Technology Office's Computational Technologies Project and led by John Schnase, Lead of CISTO's Information Science & Technology Research area, and

Tamarisk – or "saltcedar" – is an invasive shrub that was introduced into the US over a century ago from the Mediterranean. It can often outcompete native plants for water and even dry up a water source. (Image credit: USGS.)

Tom Stohlgren (USGS), the research team has significantly enhanced tools that help locate invasive species and predict their distribution. Before the ISFS, computing regional-scale models would take days, and national-scale maps were not even attempted. Now, *national*-scale maps can be generated in minutes, a fraction of the time that smaller regional maps originally took to process. Further research is being done to expand the system's capability to predict plant spread in the near-term and longer-term, and extend the project's focus to include animals and pathogens.

Both NASA and USGS are members of the National Invasive Species Council (NISC). NISC is an interdepartmental council with 13 cabinet-level member organizations formed by Executive Order in 1999 to facilitate coordination and provide leadership for federal agencies working on invasive species issues.

http://InvasiveSpecies.gsfc.nasa.gov

NASA Chairs International Direct Readout Meeting

Patrick Coronado of CISTO traveled to Benevento, Italy, October 3-6 to chair the International Earth Observing System (EOS)/National Polar-orbiting Operational Environmental Satellite System (NPOESS)/Preparatory Project (NPP) Direct Readout Meeting 2005. Hosted by the Mediterranean Agency for Remote Sensing and Environmental Control (MARSec), the meeting was attended by 101 participants from 19 countries, with 54% of the participants representing government agencies, 33% representing universities, and 13% representing the vendor community.

Direct Readout (DR) is the process of acquiring freely transmitted live satellite data. As DR technologies have become more affordable and accessible (such as with the onset of the Internet), tools have been developed by the remote sensing community to make satellite data easier to acquire, process, and utilize. As a leading member of this community, NASA supplies many of these tools to foster global data exchange and scientific collaboration. Live local and regional environmental data, in turn, benefit environmental, commercial, and public interest decision-making. DR data are increasingly employed worldwide to better understand environmental and meteorological events that affect, and at times threaten, all of us.

In the last two years the DR community has made great strides in obtaining, developing, and utilizing science algorithms for real-time and temporal applications. At this year's meeting, participants presented and discussed the latest application algorithms and systems and their impact on science, commerce, and decision-making infrastructures. Participants also discussed some of the numerous new DR tools, tutorials, and higher-level data products that are now available online.

Topics included the following:

- The US Department of Agriculture Forest Service supplies Moderate Resolution Imaging Spectroradiometer (MODIS) fire products from Terra and Aqua to firefighters to help determine resource allocation to fight wildfires in Montana.

- Research in DR polar wind data is making it possible for the National Oceanic and Atmospheric Administration to factor polar wind into hurricane forecasts in real time and improve hurricane track forecasting by 50 nautical miles.

- DR data are employed to detect fog in China to generate road hazard alerts.

- Plymouth Marine Laboratory in the United Kingdom makes MODIS ocean products available in real time to guide research ships at sea and assists the local fishing industry by detecting harmful algal blooms, locating concentrations of phytoplankton, and monitoring water quality.

- The University of Wisconsin supplies real-time tornado tracking data to the National Weather Service and the local media, and its freely distributed International MODIS/Atmospheric Infrared Sounder (AIRS) Processing Package allows DR users to calibrate, geolocate, and create environmental products from raw data.

All in all, there is a significant increase in data distribution vehicles and sources, with spatial coverage over 70% of our planet.

The last day of the meeting was largely devoted to presentations on upcoming continuity missions. Government agencies and corresponding contractors detailed the planned transitions from EOS to NPP and, ultimately, NPOESS. Highlights included presentations on science objectives and transition from MODIS to the Visible Infrared Imager/Radiometer Suite and insights into

Antenna Tower at MARSec in Benevento, Italy. (Image credit: University of Wisconsin.)

the Cross-track Infrared Sounder and the Advanced Technology Microwave Sounder. Also presented were plans for Sensor Data Record and algorithm wrapper developments at Goddard Space Flight Center's Direct Readout Laboratory and discussions of how the DR community can contribute to NPOESS calibration and validation efforts.

http://dbmeeting.gsfc.nasa.gov.
http://directreadout.gsfc.nasa.gov